

Bayesian Additive Regression Trees With Parametric Models of Heteroskedasticity

JUSTIN BLEICH^{*†} AND ADAM KAPELNER^{*‡}

The Wharton School of the University of Pennsylvania

February 24, 2014

Abstract

We incorporate heteroskedasticity into Bayesian Additive Regression Trees (**BART**) by modeling the log of the error variance parameter as a linear function of prespecified covariates. Under this scheme, the Gibbs sampling procedure for the original sum-of-trees model is easily modified, and the parameters for the variance model are updated via a Metropolis-Hastings step. We demonstrate the promise of our approach by providing more appropriate posterior predictive intervals than homoskedastic **BART** in heteroskedastic settings and demonstrating the model's resistance to overfitting. Our implementation will be offered in an upcoming release of the **R** package **bartMachine**.

1 Introduction

We consider the the following general heteroskedastic regression framework to characterize the relationship between a continuous response vector \mathbf{y} and a set of p predictor variables $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_p]$ which can be continuous or categorical:

$$\mathbf{y} = f(\mathbf{X}) + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}_n(\mathbf{0}, \sigma^2 \mathbf{D})$$

\mathbf{D} denotes the diagonal matrix whose entries are scaling factors for the error variance for each observation. In this model, the response is considered an unknown function f of the predictors and the observations, while independent, exhibit non-constant error variance. The goal of this article is to model the relationship between the predictors and response with the aim of generating accurate predictions. To this end we model f with Bayesian Additive Regression Trees (**BART**, Chipman et al., 2010) which is a sum-of-trees model that

^{*}Both authors contributed equally to this work.

[†]Electronic address: jbleich@wharton.upenn.edu; Principal Corresponding author

[‡]Electronic address: kapelner@wharton.upenn.edu; Corresponding author

has demonstrated predictive performance competitive with the best statistical learning algorithms. The original BART model is constrained to homoskedastic error variance ($\mathbf{D} = \mathbf{I}$). Here, we extend the model to flexibly handle an error variance structure which is a linear model of prespecified covariates and we name our procedure “heteroskedastic BART” or “HBART.” Similar to Huber-White sandwich estimation (White, 1980), appropriately modeling the diagonal entries of \mathbf{D} “downweights” high variance observations. This allows for (a) a more accurate model as measured by predictive performance on future observations as well as (b) posterior credible and predictive intervals which appropriately reflect the changing heteroskedasticity in predictor space.

In Section 2, we provide an overview of the literature on heteroskedastic regression modeling in a Bayesian paradigm. In Section 3, we introduce HBART, highlighting the necessary modifications to the original homoskedastic BART. In Section 4, we provide simulations to showcase the desirable properties of HBART, including less overfitting for high noise observations as well as more appropriate uncertainty intervals for predictions in the presence of heteroskedasticity. Section 5 explores two applications to real data. We conclude and offer future research directions in Section 6. The method developed in this paper will be implemented in an upcoming release of the R package `bartMachine` (Kapelner and Bleich, 2013), which is available on CRAN.

2 Bayesian Heteroskedastic Regression

Early approaches for heteroskedastic regression primarily focused on point estimation for the parameters governing the underlying heteroskedasticity of the model (for an overview, see Carroll and Ruppert, 1988). Potential problems with point estimation gave rise to a proposal of a fully Bayesian approach for heteroskedastic linear regression, where the non-constant variance depends on simple functions of an unknown parameter θ and a set of weights w_i (Boscardin and Gelman, 1994).

Cepeda and Gamerman (2001) introduce a Bayesian regression model where the conditional mean of the response is modeled as a linear function of covariates $\mathbf{x}_1, \dots, \mathbf{x}_p$ plus heteroskedastic noise. They model the variance for each observation as a monotonic differentiable function of a linear combination of another set of covariates, $g(\mathbf{z}_1, \dots, \mathbf{z}_k)$. Additionally, the function g is chosen to ensure positivity of the variance terms. The authors rely on a block Gibbs sampling approach (Geman and Geman, 1984), sampling the parameters for the mean function and variance function in two stages. In particular, the parameters for the variance function are updated via a Metropolis-Hastings step (Hastings, 1970) using the approach of Gamerman (1997), which relies on an iteratively reweighted least squares model to generate suitable proposal distributions.

More recent approaches have focused on relaxing the assumptions of linear additive components for modeling the mean and variance functions. Yau and Kohn (2003) propose nonparametric models for each of these two functions by employing penalized regression spline estimation for both models. Chan et al. (2006) extend this nonparametric model to allow for semiparametric modeling of both the mean and variance functions, using radial basis functions for nonparametric components. Additionally, their approach can handle a large number of basis terms by introducing Bayesian variable selection priors, allow-

ing model estimation to be locally adaptive. Leslie et al. (2007) relax the assumption of normal errors and developed a heteroskedastic linear regression model with general error distributions by relying on a Dirichlet process mixture prior.

Both Chan et al. (2006) and Leslie et al. (2007) rely on the sampling scheme developed in Gamerman (1997) to obtain draws from the posterior distribution of the parameters for the variance function. Our work similarly draws heavily on this technique.

3 Augmenting BART to Incorporate Heteroskedasticity

In its original formulation in Chipman et al. (2010), the authors assume that the response \mathbf{y} could be modeled as a sum-of-trees model of the covariates $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_p]$ plus homoskedastic normal noise:

$$\mathbf{Y} = \sum_{i=1}^m \mathfrak{T}_i^{\mathcal{L}}(\mathbf{X}) + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}_n(\mathbf{0}, \sigma^2 \mathbf{I}_n) \quad (1)$$

Each of the m distinct binary regression trees consists of a tree structure, denoted by \mathfrak{T} , and a set of parameters in the terminal nodes of the tree (also called “leaves”), denoted by \mathcal{L} . A tree with both its structure and set of leaf parameters is denoted by $\mathfrak{T}^{\mathcal{L}}$.

The structure of a given tree \mathfrak{T}_t includes information on how any observation recurses down the tree until a leaf node is reached. Each nonterminal, internal node contains a “splitting rule” $\mathbf{x}_j < c$. When this condition is satisfied for a given observation, the observation moves to the left daughter node or moves to the right daughter node otherwise. Once an observation lands in a terminal node, the leaf value in that node is assigned as the predicted value for that observation. $\mathcal{L}_t = \{\mu_{t,1}, \mu_{t,2}, \dots, \mu_{t,b_t}\}$ represents the set of leaf parameters for a given tree, where b_t is the number of terminal nodes in that tree.

Thus, BART estimates the mean function f using a sum-of-regression trees. Regression trees rely on recursive binary partitioning of predictor space into a menagerie of hyper-rectangles in order to approximate the unknown function of interest. By employing a sum-of-trees approach, BART is able to take advantage of the regression trees’ ability to successfully model nonlinearities and interactions while also capturing additive components of the fit. As a Bayesian model, BART is composed of a set of priors and a likelihood. Using Gibbs sampling, posterior inference for the unknown f can be obtained.

We propose an extension to BART by allowing each $\sigma_1^2, \dots, \sigma_n^2$ to be scaled by the exponential of a linear parametric function of k covariates $\mathbf{Z} := [\mathbf{z}_1, \dots, \mathbf{z}_k]$, the “heteroskedasticity covariates” which are potentially distinct from the covariates used to model the mean function, $\mathbf{x}_1, \dots, \mathbf{x}_p$. Our heteroskedastic model, HBART, is given as

$$\mathbf{y} = \sum_{i=1}^m \mathfrak{T}_i^{\mathcal{L}}(\mathbf{X}) + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}_n \left(\mathbf{0}, \sigma^2 \begin{bmatrix} \exp(\mathbf{z}_1 \cdot \boldsymbol{\gamma}) & & 0 \\ & \ddots & \\ 0 & & \exp(\mathbf{z}_n \cdot \boldsymbol{\gamma}) \end{bmatrix} \right) \quad (2)$$

where $\boldsymbol{\gamma} := [\gamma_1, \dots, \gamma_k]^\top$ is a column vector of linear coefficients for the k heteroskedasticity covariates. Thus, the variance of each observation is specified as a log-linear model:

$$\ln(\sigma_i^2) = \ln(\sigma^2) + \mathbf{z}_i \cdot \boldsymbol{\gamma} \quad \text{for } i = 1, \dots, n. \quad (3)$$

It is important to note that **BART**, by design, is an overparameterized model with “an abundance of unidentified parameters” (Chipman et al., 2010) allowing for a highly flexible fit. First, given the unidentifiable nature of the model, our focus is not on valid inference for $\boldsymbol{\gamma}$. Instead, we incorporate heteroskedasticity to aid in forecasting and generating posterior uncertainty intervals. Second, due to the already complex nature of the original **BART** algorithm, we employ parametric models for heteroskedasticity versus more sophisticated alternatives (such as the proposal of Chan et al., 2006) in order to prevent the model from becoming “too flexible.” Given highly flexible, unidentifiable estimation of both the mean and variance functions, the algorithm may have difficulty distinguishing between signal and noise, thereby shirking on its primary duty which is accurate estimation of the mean model. Hence, parametric models of heteroskedasticity represent the first step towards understanding how flexible **BART** can be in nonparametric function estimation when the homoskedasticity assumption is relaxed.

The remainder of the section is dedicated to describing the priors on **HBART** as well as the Gibbs sampling procedure for obtaining posterior inference.

3.1 Priors and Likelihood

The **BART** model requires three priors. The first prior is on the tree structures themselves and the second is on the leaf parameters. The third prior is on the error variance σ^2 . **HBART** requires these same three priors as well as a prior on $\boldsymbol{\gamma}$. By assumption, the priors on σ^2 , $\boldsymbol{\gamma}$, \mathfrak{T} and \mathcal{L} are independent of one another:

$$\begin{aligned} \mathbb{P}(\mathfrak{T}_1^{\mathcal{L}}, \dots, \mathfrak{T}_m^{\mathcal{L}}, \sigma^2, \boldsymbol{\gamma}) &= \left[\prod_t \mathbb{P}(\mathfrak{T}_t^{\mathcal{L}}) \right] \mathbb{P}(\sigma^2) \mathbb{P}(\boldsymbol{\gamma}) \\ &= \left[\prod_t \mathbb{P}(\mathcal{L}_t \mid \mathfrak{T}_t) \mathbb{P}(\mathfrak{T}_t) \right] \mathbb{P}(\sigma^2) \mathbb{P}(\boldsymbol{\gamma}) \\ &= \left[\prod_t \prod_{\ell} \mathbb{P}(\mu_{t,\ell} \mid \mathfrak{T}_t) \mathbb{P}(\mathfrak{T}_t) \right] \mathbb{P}(\sigma^2) \mathbb{P}(\boldsymbol{\gamma}) \end{aligned}$$

where the last line follows from an assumption of conditional independence of the leaf parameters given the tree structure.

The prior on the tree structures and leaf parameters are designed to provide regularization to the model by respectively preventing the trees from growing too deep and shrinking the predicted values towards the range center of the response. These priors, as well as a prior on proposing splitting rules, are the same as those used in the original **BART** model and we refer the interested reader to Chipman et al. (2010).

For the homoskedastic implementation, the prior is on the error variance and is chosen to be $\sigma^2 \sim \text{InvGamma}(\nu/2, \nu\lambda/2)$. λ is determined from the data so that there is a

$q = 90\%$ a priori chance (by default) that the **BART** model will improve upon the root mean square error (RMSE) from an ordinary least squares regression (therefore, the majority of the prior probability mass lies below the RMSE of a least squares regression). Additionally, this prior limits the probability mass placed on small values of σ^2 to prevent overfitting. We use this same data-informed prior for the σ^2 parameter in **HBART**, the logarithm of which serves as the intercept term in the log-linear model for the variances (Equation 3).

Following the approach of Gamerman (1997), we place a multivariate normal prior on γ given as:

$$\gamma \sim \mathcal{N}_k \left(\gamma_0, \begin{bmatrix} \Sigma_{11} & & 0 \\ & \ddots & \\ 0 & & \Sigma_{kk} \end{bmatrix} \right) \quad (4)$$

We make the simplifying assumption that each component of the prior is independent of one another, but this can be easily generalized.

Along with a set of priors, **HBART** (and **BART**) consists of the likelihood of responses in the leaf nodes. The likelihood is assumed to be normal with the mean being the “best guess” of the leaf parameters at the current moment and variance being the best guess of the variance at the moment i.e. $\mathbf{y}_\ell \sim \mathcal{N}(\mu_\ell, \sigma_\ell^2/m)$. These “best guesses” are the values being conditioned on in the Gibbs sampler during each iteration.

3.2 Sampling from the Posterior

The Gibbs sampler can be used to obtain draws from $\mathbb{P}(\mathfrak{F}_1^{\mathcal{E}}, \dots, \mathfrak{F}_m^{\mathcal{E}}, \sigma^2, \gamma \mid \mathbf{y}, \mathbf{X})$, the posterior distribution of the model parameters. As with the original sampling scheme for **BART**, **HBART** also relies on “Bayesian backfitting” (Hastie and Tibshirani, 2000) to fit each tree iteratively, holding all other $m - 1$ trees constant. This is achieved by considering the residual response when updating the j^{th} tree $\mathbf{R}_j := \mathbf{y} - \sum_{t \neq j} \mathfrak{F}_t^{\mathcal{E}}(\mathbf{X})$.

The Gibbs sampler for **HBART** proceeds by first proposing a change to the first tree’s structure \mathfrak{F}_1 which are accepted or rejected via a Metropolis-Hastings step (Hastings, 1970). Tree structures are altered by introducing small changes: growing a terminal node by adding two terminal daughter nodes, pruning two terminal daughter nodes, or changing a split rule. We denote these possible alterations as: GROW, PRUNE, and CHANGE.¹ Given the tree structure, samples from the posterior of the b leaf parameters $\mathcal{E}_1 := \{\mu_1, \dots, \mu_b\}$ are then drawn from the conjugate-normal posterior distribution. This procedure proceeds iteratively for each tree, using an updated set of partial residuals \mathbf{R}_j .

Once each tree structure and leaf values has been updated, a draw from the posterior of σ^2 conditional on all other parameters is made based on the full model residuals $\mathcal{E} := \mathbf{y} - \sum_{t=1}^m \mathfrak{F}_t^{\mathcal{E}}(\mathbf{X})$. Finally, a draw from the posterior of γ conditional the other parameters is computed via a Metropolis-Hastings step.

The steps of the full procedure are illustrated below:

¹An additional step known as SWAP was proposed in Chipman et al. (2010), but this step is not implemented in the **bartMachine** package which was used to develop **HBART**.

$$\begin{aligned}
1: \quad \mathfrak{T}_1 & \mid R_{-1}, \sigma^2, \gamma \\
2: \quad \mathcal{L}_1 & \mid \mathfrak{T}_1, R_{-1}, \sigma^2, \gamma \\
& \vdots \\
2m-1: \quad \mathfrak{T}_m & \mid R_{-m}, \sigma^2, \gamma \\
2m: \quad \mathcal{L}_m & \mid \mathfrak{T}_m, R_{-m}, \sigma^2, \gamma \\
2m+1: \quad \sigma^2 & \mid \mathfrak{T}_1^{\mathcal{L}}, \dots, \mathfrak{T}_m^{\mathcal{L}}, \gamma, \mathcal{E} \\
2m+2: \quad \gamma & \mid \sigma^2, \mathcal{E}
\end{aligned} \tag{5}$$

All $2m + 2$ steps represent a *single* Gibbs iteration². After a sufficient burn-in period, burned-in draws from the posterior of f are obtained. A point prediction $\hat{f}(\mathbf{x})$ can be obtained by taking the average of the burned-in values of $\mathfrak{T}_1^{\mathcal{L}}, \dots, \mathfrak{T}_m^{\mathcal{L}}$ evaluated at a given \mathbf{x} ; posterior credible intervals for f are computed by using the quantiles of the burned-in values. Posterior predictive intervals at a given \mathbf{x} are computed as follows:

1. Draw $f(\mathbf{x})$ via
 - i) drawing one of the burned-in sum-of-trees collections $\mathfrak{T}_1^{\mathcal{L}}, \dots, \mathfrak{T}_m^{\mathcal{L}}$ and
 - ii) computing $f(\mathbf{x}) = \sum_{i=1}^m \mathfrak{T}_i^{\mathcal{L}}(\mathbf{x})$.
2. Draw $\sigma^2(\mathbf{x})$ via
 - i) obtaining γ and σ^2 from the same Gibbs sample from which the sum-of-trees was obtained,
 - ii) determining the \mathbf{z} which corresponds to the \mathbf{x} of interest and
 - iii) computing $\sigma^2(\mathbf{x})$ by evaluating γ , σ^2 and \mathbf{z} in the exponentiation of Equation 3.
3. Take one draw from $\mathcal{N}(f(\mathbf{x}), \sigma^2(\mathbf{x}))$ which is the BART model (Equation 1)
4. Collect draws from step 3 by repeating steps 1–3 many times. Then, return the desired quantiles.

Note that HBART requires modifications to the original BART likelihood calculations necessary for Metropolis-Hastings steps to alter the tree structures (steps 1, 3, \dots , $2m - 1$ of Equation 5). Also, the posterior distributions for the leaf parameters must be updated (steps 2, 4, \dots , $2m$ of Equation 5). It is worth noting that these modifications are valid for any heteroskedastic BART model and not just the log-linear HBART model proposed in this work, as they computed as functions of $\sigma_1^2, \dots, \sigma_n^2$. Additionally, modifications to the posterior distribution of σ^2 (step $2m + 1$ of Equation 5) is required. Finally, the sample of γ (step $2m + 2$ of Equation 5) is obtained using the Metropolis-Hastings procedure with the proposal distribution outlined in Gamerman (1997). We provide explicit computational details for each of these steps in Appendix A.

²BART relies on a similar scheme, removing the conditioning on γ at each step and not requiring step $2m + 2$.

4 Simulations

4.1 Univariate Model

We begin by comparing the performance of **BART** versus **HBART** in a simple univariate setting. Consider a single predictor \mathbf{x} which is a uniformly spaced sequence of n points on $[0, 1]$. Then, we consider two models. The first is heteroskedastic and is given by:

$$Y_i = 100x_i + \mathcal{E}_i, \quad \mathcal{E}_i \stackrel{\text{ind}}{\sim} \mathcal{N}(0, \sigma_i^2), \quad \sigma_i^2 = \exp(7x_i) \quad (6)$$

The second is homoskedastic and is given by:

$$Y_i = 100x_i + \mathcal{E}_i, \quad \mathcal{E}_i \stackrel{\text{ind}}{\sim} \mathcal{N}(0, 5^2) \quad (7)$$


For **HBART**, the matrix of covariates for the parametric variance model will be set to $\mathbf{Z} = [\mathbf{x}]$, our one uniformly spaced covariate. Parenthetically, note that the default in our software implementation is to set $\mathbf{Z} = \mathbf{X}$. Therefore, **HBART** will have a correctly specified variance function for the model given in Equation 6. We include the homoskedastic model as well to benchmark **HBART**'s performance to determine if the unnecessary extra complexity of the variance model degrades the algorithm's performance.

Figure 1 compares **BART** to **HBART** in both the heteroskedastic and homoskedastic simulated models by gauging two metrics: accuracy of estimation of f and appropriateness of posterior predictive intervals.


Figure 1a highlights the posterior means \hat{f} estimated by the two different models. Note that **HBART** provides a better estimate of the true f when $x > 0.7$, the region of relatively high variance. By estimating large variance in this region, it has the flexibility to down-weight these high variance observations, allowing for more shrinkage towards the global average and away from the noisy local (within-node) sample mean. In contrast, **BART** overfits in this region. By assuming homoskedasticity, the algorithm is handicapped, and is obligated to move its mean function up when the noise term is large and positive and down when the noise term is large and negative. Note that both algorithms perform well when $x < 0.7$, where the data has lower variance.

Figure 1b provides 90% posterior predictive intervals for future observations (as explained in the procedure outlined in Section 3.2). Given the homoskedasticity assumption of **BART**, the prediction intervals at each x -location are of constant width. This implies that the intervals are too wide at low values of x and too narrow at higher values of x . In contrast, **HBART** provides more appropriate prediction intervals, narrow at low x and wide at high x , thus correctly reflecting the heteroskedasticity in the underlying data-generating model. Although not the primary focus of this paper, examining the burned in γ values yielded a 90% credible interval of $[6.52, 7.65]$, which captures 7, the value of the linear coefficient in the log-linear variance model of Equation 6.


For the homoskedastic model, Figures 1c and 1d highlight that **BART** and **HBART** yield virtually identical results in terms of mean function estimation and predictive intervals. Thus, **HBART** seems to be robust in the absence of heteroskedasticity for this illustration.




(a) Heteroskedastic Model,
Posterior Mean Estimates



(b) Heteroskedastic Model,
Posterior Predictive Intervals



(c) Homoskedastic Model,
Posterior Mean Estimates



(d) Homoskedastic Model,
Posterior Predictive Intervals

Figure 1: BART's and HBART's posterior mean estimates and 90% posterior predictive intervals for each algorithm built from a sample of $n = 250$ observations drawn from the processes in Equation 6 (a and b) and Equation 7 (c and d). **Red** lines correspond to the results of a BART model and **blue** lines correspond to the results from a HBART model. The black lines in (b) and (d) correspond to the true conditional mean function ($Y_i = 100x_i$) and the black '+'s represent the actual observations.

4.2 Multivariate Model

We now consider the following data generating process, similar to the model simulated in Cepeda and Gamerman (2001):

$$Y_i = f(\mathbf{x}_i) + \mathcal{E}_i, \quad f(\mathbf{x}_i) = -35 + .35x_{1,i} - 1.7x_{2,i}, \quad \mathcal{E}_i \stackrel{iid}{\sim} \mathcal{N}(0, \sigma_i^2) \quad (8)$$

$$x_{1,i} \stackrel{iid}{\sim} \text{U}(0, 400), \quad x_{2,i} \stackrel{iid}{\sim} \text{U}(10, 23), \quad x_{3,i} \stackrel{iid}{\sim} \text{U}(0, 10) \quad (9)$$

$$\sigma_i^2 = \exp(-6 + .03x_{1,i} + .4x_{3,i})$$

We set the number of observations to be $n = 500$ and then we fit an HBART model using the default $\mathbf{Z} = [\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3]$. Thus, the variance model is misspecified (the true model is not a function of \mathbf{x}_2). We again compare HBART's performance to that of BART.




Figure 2: Estimates of the conditional mean function f for HBART and BART with associated 90% credible intervals. The x-axis is the true value of the conditional mean function and the y-value is the model estimate. Gray lines illustrate 90% credible intervals. If the true conditional mean falls within the interval, the point is colored **green** and points in **red** signify the true value falls outside of the interval.

Figure 2 plots actual values of the conditional mean function versus fitted values. This illustration demonstrates that in areas of high variance, BART has difficulty separating the mean function from the noise, and as a result, provides wide credible intervals for the true f . HBART, on the other hand, provides more narrow credible intervals, indicating that the

algorithm was able to successfully separate the mean function from the heteroskedastic variance structure in the data generating process.

We next evaluate the performance of **HBART** versus **BART** in terms of out-of-sample predictive performance. We consider two models: one with homoskedastic errors, $\mathcal{E} \stackrel{iid}{\sim} \mathcal{N}(0, 3^2)$, and one with heteroskedastic error structure according to Equation 9. For each error structure, we then generate $n = 500$ training observations based on the data generating process outlined in Equation 8. Both **HBART** and **BART** models are constructed on the training set and performance is evaluated in terms of root mean square error (RMSE) on an additional $n = 500$ independent test observations drawn from the same data generating process.

Figure 3 displays the out-of-sample performance results for 100 simulations. For the heteroskedastic model, **HBART** significantly outperforms **BART** (Figure 3a). For the homoskedastic model, the performance of **BART** and **HBART** are statistically equal and the results of both algorithms are quite stable (Figure 3b). This plot provides more evidence of the robustness of **HBART**'s performance in the absence of heteroskedasticity.

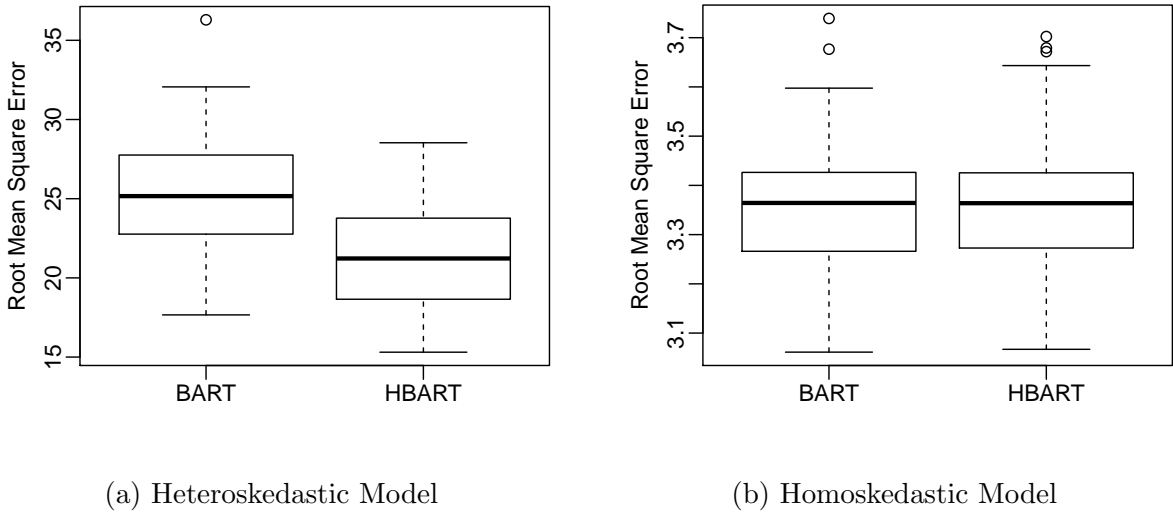


Figure 3: Distribution of out-of-sample RMSE for **BART** and **HBART** fit in 100 simulations for (a) the heteroskedastic data generating process of Equations 8 and 9 and (b) the homoskedastic data generating process of Equation 8 with $\mathcal{E} \stackrel{iid}{\sim} \mathcal{N}(0, 3^2)$. The models are built on 500 observations and tested on 500 independent observations.

5 Real Data Examples

5.1 Lidar Data

We consider the LIDAR data set explored in Ruppert et al. (2003). The data set contains 221 observations. The response denoted **Log(Ratio)** is the logarithm of the ratio of reflected laser-emitted light from two sources and the predictor denoted **Range** is the distance traveled before the light is reflected back to its source. Leslie et al. (2007) explored this data set by

fitting both nonparametric mean and variance functions to the data under the assumptions of both normal and non-normal error distributions.

We fit the data using BART and HBART. For HBART, \mathbf{Z} is taken to be $[\text{Range}, \text{Range}^2]$ where the two columns are orthogonalized. Figure 4 illustrates the posterior mean estimates for both HBART and BART. One will notice that both algorithms estimate relatively similar posterior mean functions with HBART’s estimation being slightly more smooth in the region of higher variance than that of BART. Figure 1b shows 90% posterior predictive intervals for the two algorithms. The intervals for HBART seem to appropriately reflect the heteroskedasticity in the data, while BART’s predictive intervals are too wide at low values of *Range* and too narrow at high values of *Range*.

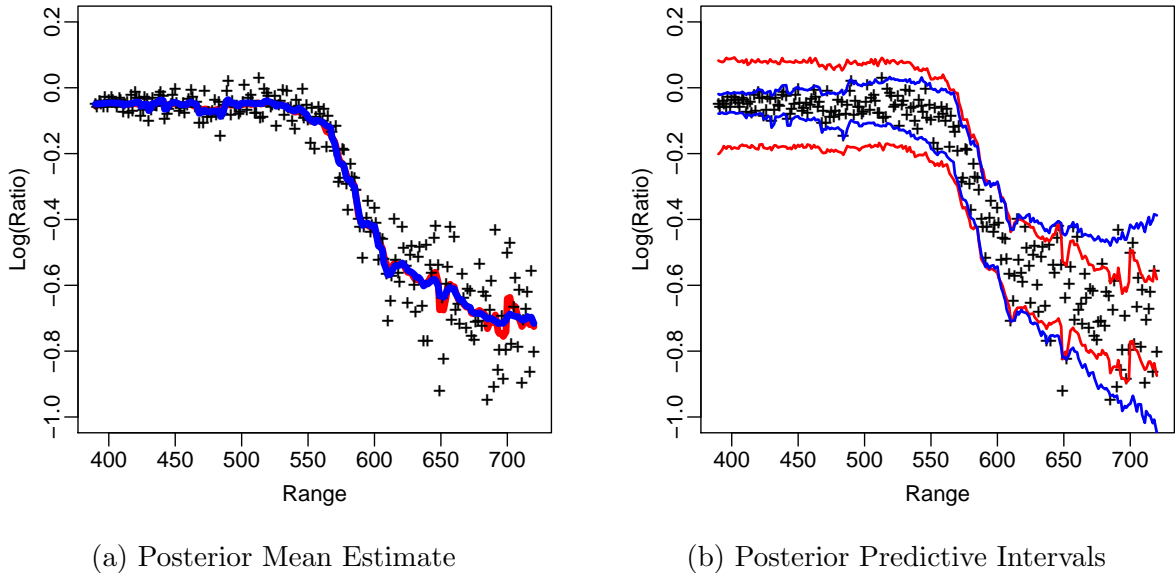


Figure 4: BART’s and HBART’s posterior mean estimates and 90% posterior predictive intervals for the Lidar data. **Red** lines correspond to the BART model and **blue** lines correspond to the HBART model.

5.2 Motorcycle Data

We next consider a dataset of simulated motorcycle crashes that was compiled by Schmidt et al. (1981). The observations in the data set consist of accelerometer readings (*acceleration*) taken from riders’ helmets at 133 different points in time (*time*) after a simulated impact. As discussed in Gramacy (2007), many researchers find that this dataset exhibits multiple regimes in both the mean function and variance function over time.

This dataset was also explored by Taddy et al. (2011) who remarked that the 90% posterior predictive interval for BART appeared to “variously over or under estimate data uncertainty around the regression mean. In particular, BART’s global variance term is misspecified for this heteroskedastic data.” We attempt to remedy this problem with HBART. Exploring a scatterplot of the data, the model seems to be characterized by a low variance regime followed by a high variance regime and then an additional low variance

regime (see Figure 5). Hence, when building an **HBART** model, we do not use the default \mathbf{Z} . To capture the low-high-low variance relationship, we specify the model quadratic in the predictor, $\mathbf{Z} = [\text{time}, \text{time}^2]$ where the two columns are orthogonalized.

Figure 5 displays 90% posterior predictive intervals for **HBART**, **BART**, dynamic trees (**dynaTree**, Taddy et al., 2011), and treed Gaussian processes (**tgp**, Gramacy, 2007). Each of the algorithms has a similar estimate of the posterior mean process (unshown), but there are some differences in the posterior predictive intervals. **BART**’s predictive intervals are much too wide at low and high values of time and perhaps too narrow for the intermediate values of time. However, **HBART**’s intervals are quite similar to those of **dynaTree**, being widest for intermediate values of time and more narrow near the beginning and end. Interestingly, **HBART** is the only of the four models that builds a narrower predictive interval at the higher values of time.




Figure 5: 90% posterior predictive intervals for **BART**, **HBART**, **dynaTree** and **TGP**.

6 Discussion

We have proposed **HBART**, an extension of **BART** that relaxes the assumption of homoskedasticity in the model errors. In particular, we have developed a model that allows for a multiplicative heteroskedastic error structure, where the multiplicative factor is the exponential of a parametric function of some set of covariates.

Through simulations and explorations of real data, we have demonstrated **HBART**’s potential for generating more appropriate posterior predictive intervals in the presence of appropriately modeled heteroskedastic data versus **BART**. Additionally, **BART** suffers from overfitting in areas of very high variance and **HBART** seems to offer promise in ameliorating this issue. In our explorations, the added complexity of fitting a model to the error terms

did not hinder HBART’s performance on homoskedastic data. In this situation, HBART’s estimates of the posterior means, predictive intervals, and out-of-sample RMSE were very similar to BART’s.

Originally proposed in Chipman et al. (2010) and implemented in Kapelner and Bleich (2013), it is possible to cross-validate over the a number of the hyperparameters of the BART model. Future work will involve extending BART-CV to HBART-CV, where it is possible to cross-validate over a selection of prior variances for the parametric variance model to impose varying degrees of shrinkage by modifying the Σ_{jj} hyperparameters. One final direction of research is to further relax the assumptions on the error structure. For instance, one could incorporate more flexible variance models such as smoothing splines instead of the standard linear model or relax the assumption of normality of the errors by considering Dirichlet mixture priors.

Replication

The code for HBART is located at <http://github.com/kapelner/bartMachine> in the branch “hBART” and it is open source under the MIT license. The results, tables, and figures found in this paper can be replicated via the `replication.R` script located in the `hbart_paper` folder within this git repository. The code can be conveniently installed as a local R package.

Acknowledgements

We thank Richard Berk, Ed George and Abba Krieger for helpful discussions. We thank Simon Urbanek for his generous help with rJava. Adam Kapelner acknowledges support from the National Science Foundation’s Graduate Research Fellowship Program.

References

- Boscardin, W. and Gelman, A. (1994). Bayesian Computation for Parametric Models of Heteroscedasticity in the Linear Model. *University of California, Berkeley*.
- Carroll, R. and Ruppert, D. (1988). *Transformation and Weighting in Regression*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis.
- Cepeda, E. and Gamerman, D. (2001). Bayesian modeling of variance heterogeneity in normal regression models. *Brazilian Journal of Probability and Statistics*, pages 207–221.
- Chan, D., Kohn, R., Nott, D., and Kirby, C. (2006). Locally adaptive semiparametric estimation of the mean and variance functions in regression models. *Journal of Computational and Graphical Statistics*, pages 0–26.
- Chipman, H., George, E., and McCulloch, R. (2010). BART: Bayesian Additive Regressive Trees. *The Annals of Applied Statistics*, 4(1):266–298.
- Gamerman, D. (1997). Sampling from the posterior distribution in generalized linear mixed models. *Statistics and Computing*.

- Gelman, A., Carlin, J., Stern, H., and Rubin, D. (2004). *Bayesian Data Analysis*. Chapman & Hall / CRC, second edition.
- Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 6:721–741.
- Gramacy, R. B. (2007). tgp: An R Package for Bayesian Nonstationary, Semiparametric Nonlinear Regression and Design by Treed Gaussian Process Models. *Journal Of Statistical Software*, 19(9).
- Hastie, T. and Tibshirani, R. (2000). Bayesian Backfitting. *Statistical Science*, 15(3):196–213.
- Hastings, W. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109.
- Kapelner, A. and Bleich, J. (2013). bartmachine: Machine learning with bayesian additive regression trees. *ArXiv*.
- Leslie, D., Kohn, R., and Nott, D. (2007). A general approach to heteroscedastic linear regression. *Statistics and Computing*, pages 1–29.
- Ruppert, D., Wand, P., and Carroll, R. (2003). *Semiparametric Regression*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Schmidt, G., Mattern, R., and Schöler, F. (1981). Biomechanical investigation to determine physical and traumatological differentiation criteria for the maximum load capacity of head and vertebral column with and without protective helmet under effect. Technical report, EFC Research Program on Biomechanics of Impacts. Final report Phase III, Project G5. Institute für Rechtsmedizin, Universität Heidelberg.
- Taddy, M., Gramacy, R., and Polson, N. (2011). Dynamic Trees for Learning and Design. *Journal of the American Statistical Association*, 106(493):109–123.
- White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, pages 817–838.
- Yau, P. and Kohn, R. (2003). Estimation and variable selection in nonparametric heteroscedastic regression. *Statistics and Computing*, pages 191–208.

A Gibbs Sampling Details for HBART

In the following sections, we provide implementation details for each step of the HBART Gibbs sampler (Equation 5).

Drawing σ^2 (step $2m + 1$)

Drawing σ^2 in HBART requires a slight modification to the scheme used in the original homoskedastic BART. With the errors distributed multivariate normal (Equation 2) and the prior on σ^2 being distributed inverse-gamma, conjugacy yields the posterior as an inverse-gamma distribution just as in BART. The quadratic form of the errors and their covariance matrix carries into the scale parameter. Hence, we sample σ^2 as

$$\begin{aligned}\sigma^2 &\sim \text{InvGamma}\left(\frac{\nu + n}{2}, \frac{\nu\lambda + \boldsymbol{\mathcal{E}}^T \mathbf{D}^{-1} \boldsymbol{\mathcal{E}}}{2}\right) \\ &= \text{InvGamma}\left(\frac{\nu + n}{2}, \frac{\nu\lambda}{2} + \frac{1}{2} \left(\sum_{i=1}^n \exp(-\mathbf{z}_i^\top \boldsymbol{\gamma}) \mathcal{E}_i^2 \right)\right).\end{aligned}$$

The default values of hyperparameters ν and q are set to be the same as the defaults of the original homoskedastic BART algorithm.

Drawing $\boldsymbol{\gamma}$ (step $2m + 2$)

The posterior of $\boldsymbol{\gamma} \mid \mathcal{E}_1, \dots, \mathcal{E}_n, \sigma^2$ is proportional to:

$$\begin{aligned}\mathbb{P}(\boldsymbol{\gamma} \mid \mathcal{E}_1, \dots, \mathcal{E}_n, \sigma^2) &\propto \mathbb{P}(\mathcal{E}_1, \dots, \mathcal{E}_n \mid \boldsymbol{\gamma}, \sigma^2) \mathbb{P}(\boldsymbol{\gamma}) \\ &= \prod_{i=1}^n \frac{1}{\sqrt{2\pi\sigma^2 \exp(\mathbf{z}_i^\top \boldsymbol{\gamma})}} \exp\left(-\frac{1}{2\sigma^2 \exp(\mathbf{z}_i^\top \boldsymbol{\gamma})} \mathcal{E}_i^2\right) \times \\ &\quad \prod_{j=1}^p \frac{1}{\sqrt{2\pi\Sigma_{jj}}} \exp\left(-\frac{1}{2\Sigma_{jj}} (\gamma_j - \gamma_{0j})^2\right) \\ &\propto \frac{1}{\sqrt{\sigma^2 \exp(\sum_{i=1}^n \mathbf{z}_i^\top \boldsymbol{\gamma})}} \exp\left(-\frac{1}{2} \left(\sum_{i=1}^n \frac{\mathcal{E}_i^2}{\sigma^2 \exp(\mathbf{z}_i^\top \boldsymbol{\gamma})} + \sum_{j=1}^p \frac{(\gamma_j - \gamma_{0j})^2}{\Sigma_{jj}} \right)\right)\end{aligned}$$

Since it is not tractable to draw directly from the above distribution, a proposal distribution is required. A suitable proposal distribution can be obtained using Gamerman (1997), which is based on a single step of an iteratively reweighted least squares algorithm. Letting $\boldsymbol{\gamma}$ represent the current value of the parameter vector, the recommended proposal density is $\mathbb{P}(\boldsymbol{\gamma}^* \mid \boldsymbol{\gamma}) = \mathcal{N}_p(\mathbf{a}(\boldsymbol{\gamma}), \mathbf{B})$ where

$$\begin{aligned}\mathbf{a} &= \mathbf{B} \left(\boldsymbol{\Sigma}^{-1} \boldsymbol{\gamma}_0 + \frac{1}{2} \mathbf{Z}^T \mathbf{w} \right), \quad \mathbf{B} = \left(\boldsymbol{\Sigma}^{-1} + \frac{1}{2} \mathbf{Z}^T \mathbf{Z} \right)^{-1} \text{ and} \\ \mathbf{w} &= \left[\mathbf{z}_1^\top \boldsymbol{\gamma} + \frac{\mathcal{E}_1^2}{\sigma^2 \exp(\mathbf{z}_1^\top \boldsymbol{\gamma})} - 1 \dots \mathbf{z}_n^\top \boldsymbol{\gamma} + \frac{\mathcal{E}_n^2}{\sigma^2 \exp(\mathbf{z}_n^\top \boldsymbol{\gamma})} - 1 \right]^\top.\end{aligned}$$

Then, $\boldsymbol{\gamma}^*$ is accepted if a draw from a standard uniform distribution is less than the Metropolis-Hastings ratio (Gelman et al., 2004, p.291),

$$r = \underbrace{\frac{\mathbb{P}(\gamma^*|\gamma)}{\mathbb{P}(\gamma|\gamma^*)}}_{\text{jump ratio}} \underbrace{\frac{\mathbb{P}(\gamma^* | \sigma^2, \mathcal{E}_1, \dots, \mathcal{E}_n)}{\mathbb{P}(\gamma | \sigma^2, \mathcal{E}_1, \dots, \mathcal{E}_n)}}_{\text{likelihood ratio}}. \quad (10)$$

The *jump ratio* in Equation 10 becomes

$$\begin{aligned} \frac{\mathbb{P}(\gamma^*|\gamma)}{\mathbb{P}(\gamma|\gamma^*)} &= \frac{(2\pi)^{-\frac{p}{2}} |\mathbf{B}|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} (\gamma^* - a(\gamma))^\top \mathbf{B}^{-1} (\gamma^* - a(\gamma))\right)}{(2\pi)^{-\frac{p}{2}} |\mathbf{B}|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} (\gamma - a(\gamma^*))^\top \mathbf{B}^{-1} (\gamma - a(\gamma^*))\right)} \\ &= \exp\left(\frac{1}{2} (\gamma - \gamma^* + a(\gamma) - a(\gamma^*))^\top \mathbf{B}^{-1} (\gamma - \gamma^* + a(\gamma) - a(\gamma^*))\right), \end{aligned}$$

the *likelihood ratio* in Equation 10 is calculated to be

$$\begin{aligned} \frac{\mathbb{P}(\gamma^* | \sigma^2, \mathcal{E}_1, \dots, \mathcal{E}_n)}{\mathbb{P}(\gamma | \sigma^2, \mathcal{E}_1, \dots, \mathcal{E}_n)} &= \frac{\prod_{i=1}^n \frac{1}{\sqrt{2\pi\sigma^2 \exp(z_i^\top \gamma^*)}} \exp\left(-\frac{1}{2\sigma^2 \exp(z_i^\top \gamma^*)} \mathcal{E}_i^2\right)}{\prod_{i=1}^n \frac{1}{\sqrt{2\pi\sigma^2 \exp(z_i^\top \gamma)}} \exp\left(-\frac{1}{2\sigma^2 \exp(z_i^\top \gamma)} \mathcal{E}_i^2\right)} \\ &\quad \times \frac{\prod_{j=1}^p \frac{1}{\sqrt{2\pi\Sigma_{jj}}} \exp\left(-\frac{1}{2\Sigma_{jj}} (\gamma_j^* - \gamma_{0j})^2\right)}{\prod_{j=1}^p \frac{1}{\sqrt{2\pi\Sigma_{jj}}} \exp\left(-\frac{1}{2\Sigma_{jj}} (\gamma_j - \gamma_{0j})^2\right)} \\ &= \exp\left(\frac{1}{2} \left(\sum_{i=1}^n z_i^\top (\gamma - \gamma^*) \right. \right. \\ &\quad \left. \left. + \frac{1}{\sigma^2} \left(\sum_{i=1}^n \mathcal{E}_i^2 \left(\exp(-z_i^\top \gamma) - \exp(-z_i^\top \gamma^*) \right) \right) \right) \right) \\ &\quad \left. + \sum_{j=1}^p \frac{1}{\Sigma_{jj}} \left(\gamma_j^2 - \gamma_j^{*2} + 2\gamma_{0j} (\gamma_j^* - \gamma_j) \right) \right) \end{aligned}$$

which results in the Metropolis-Hastings ratio of Equation 10 of

$$\begin{aligned} r &= \exp\left(\frac{1}{2} \left((\gamma - \gamma^* + a(\gamma) - a(\gamma^*))^\top \mathbf{B}^{-1} (\gamma + \gamma^* - (a(\gamma) + a(\gamma^*))) \right. \right. \\ &\quad \left. \left. + \sum_{i=1}^n z_i^\top (\gamma - \gamma^*) + \frac{1}{\sigma^2} \sum_{i=1}^n \mathcal{E}_i^2 \left(\exp(-z_i^\top \gamma) - \exp(-z_i^\top \gamma^*) \right) \right) \right) \\ &\quad \left. + \sum_{j=1}^p \frac{1}{\Sigma_{jj}} \left(\gamma_j^2 - \gamma_j^{*2} + 2\gamma_{0j} (\gamma_j^* - \gamma_j) \right) \right). \end{aligned}$$

We now choose default values for the hyperparameters $\gamma_{0,1}, \dots, \gamma_{0,p}$ and $\Sigma_{11}, \dots, \Sigma_{pp}$. All $\gamma_{0,j}$ are set to 0. This choice centers the prior distribution of the linear model coefficients at zero. For the variance hyperparameters, we choose the Σ_{jj} 's to be 1000 for all j which is sufficiently large so our model will not impose shrinkage of the coefficients towards 0. Investigating this parameter's role in our algorithm we view as fruitful future work.

Drawing the \varnothing 's (steps 2, 4, \dots , $2m$)

Sampling the leaf parameters must be adjusted to reflect the heteroskedasticity in the model. Observations considered highly variable will now *be downweighted* when constructing an estimate of a terminal node's prediction. Recall that $\varnothing_t \mid \mathfrak{F}_1, \mathbf{R}_1, \sigma_1^2, \dots, \sigma_n^2$ is the sampling for all leaves where each leaf is considered independent,

$$\begin{array}{c|c} \mu_{t1} & \mathfrak{F}_t, \mathbf{R}_{t1}, \sigma_1^2, \dots, \sigma_n^2 \\ \mu_{t2} & \mathfrak{F}_t, \mathbf{R}_{t2}, \sigma_1^2, \dots, \sigma_n^2 \\ \vdots & \\ \mu_{tb_t} & \mathfrak{F}_t, \mathbf{R}_{tb_t}, \sigma_1^2, \dots, \sigma_n^2 \end{array}$$

where b_t denotes the number of terminal nodes in the t th tree and the subscripts on the R_{t_i} terms denote only the data that is apportioned to the specific terminal node. Recall that the prior on the leaf parameters, μ 's, are normal and the likelihood of the responses, R_t 's, are assumed normal as well.

Given the normal-normal conjugacy, we derive the normal posterior distribution for a given leaf's prediction which is necessarily a function of the heterogeneous variances. We drop the subscripts on the R term for convenience and denote k as the number of data records that fell into the given leaf. Note that we drop double subscripting on the variances $\{\sigma_1^2, \dots, \sigma_k^2\}$ which are a subset of $\{\sigma_1^2, \dots, \sigma_n^2\}$ for the data records in this leaf.

$$\begin{aligned} \mathbb{P}(\mu \mid \mathbf{R}, \sigma_1^2, \dots, \sigma_k^2, \gamma) &\propto \mathbb{P}(R \mid \mu, \sigma_1^2, \dots, \sigma_k^2, \gamma) \mathbb{P}(\mu \mid \sigma_1^2, \dots, \sigma_k^2) \\ &= \mathcal{N}_k(\mu \mathbf{1}, \sigma^2 \mathbf{D}) \mathcal{N}(0, \sigma_\mu^2) = \left(\prod_{i=1}^k \mathcal{N}(\mu, \sigma_i^2) \right) \mathcal{N}(0, \sigma_\mu^2) \\ &= \mathcal{N} \left(\frac{\sum_{i=1}^k \frac{R_i}{\sigma_i^2}}{\frac{1}{\sigma_\mu^2} + \sum_{i=1}^k \frac{1}{\sigma_i^2}}, \frac{1}{\frac{1}{\sigma_\mu^2} + \sum_{i=1}^k \frac{1}{\sigma_i^2}} \right) \end{aligned}$$

Drawing the \mathfrak{F} 's (steps 1, 3, \dots , $2m - 1$)

As described in Kapelner and Bleich (2013, Appendix A), the draw of a new tree structure relies on a Metropolis-Hastings step where trees can be altered via growing new daughter

nodes from an existing terminal node (**GROW**), pruning two terminal nodes such that their parent becomes terminal (**PRUNE**), or changing the splitting rule in a node (**CHANGE**).

Below is the Metropolis-Hastings ratio where the parameter sampled is the tree and the data is the responses unexplained by other trees denoted by \mathbf{R} . We denote the new, proposal tree with an asterisk and the original tree without the asterisk.

$$r = \frac{\mathbb{P}(\mathfrak{T}_* \rightarrow \mathfrak{T}) \mathbb{P}(\mathfrak{T}_* | \mathbf{R}, \sigma^2, \gamma)}{\mathbb{P}(\mathfrak{T} \rightarrow \mathfrak{T}_*) \mathbb{P}(\mathfrak{T} | \mathbf{R}, \sigma^2, \gamma)} \quad (11)$$

This can be reformulated using repeated applications of Bayes Rule to be a product of three ratios.

$$r = \underbrace{\frac{\mathbb{P}(\mathfrak{T}_* \rightarrow \mathfrak{T})}{\mathbb{P}(\mathfrak{T} \rightarrow \mathfrak{T}_*)}}_{\text{transition ratio}} \times \underbrace{\frac{\mathbb{P}(\mathbf{R} | \mathfrak{T}_*, \sigma^2, \gamma)}{\mathbb{P}(\mathbf{R} | \mathfrak{T}, \sigma^2, \gamma)}}_{\text{likelihood ratio}} \times \underbrace{\frac{\mathbb{P}(\mathfrak{T}_*)}{\mathbb{P}(\mathfrak{T})}}_{\text{tree structure ratio}}$$

Note that the probability of the tree structure is independent of σ^2 and γ .

The transition ratio and the tree structure ratio remain the same as in the original **BART** as they do not depend on the variance parameters. The likelihood ratio now must take into account the heterogeneity in variances. The **PRUNE** likelihood ratio is the inverse of the **GROW** likelihood ratio. Thus, we only need to compute likelihood ratios for **GROW** and **CHANGE**.

The Likelihood Ratio for the **GROW** proposal

Since the likelihoods are solely determined by the terminal nodes, the grown proposal tree differs from the original tree by only the selected node to be grown. We denote the node to be grown by ℓ , the left child by ℓ_L and the right child by ℓ_R .

$$\frac{\mathbb{P}(\mathbf{R} | \mathfrak{T}_*, \sigma^2, \gamma)}{\mathbb{P}(\mathbf{R} | \mathfrak{T}, \sigma^2, \gamma)} = \frac{\int_{\mathbb{R}} \mathbb{P}(\mathbf{R}_{\ell_L} | \mu, \sigma^2, \gamma) \mathbb{P}(\mu) d\mu \int_{\mathbb{R}} \mathbb{P}(\mathbf{R}_{\ell_R} | \mu, \sigma^2, \gamma) \mathbb{P}(\mu) d\mu}{\int_{\mathbb{R}} \mathbb{P}(\mathbf{R}_{\ell} | \mu, \sigma^2, \gamma) \mathbb{P}(\mu) d\mu} \quad (12)$$

Each of these three integrals are the same with regards to marginalizing the μ term:

$$\begin{aligned} & \int_{\mathbb{R}} \mathbb{P}(\mathbf{R} | \mu, \sigma^2, \gamma) \mathbb{P}(\mu) d\mu \\ &= \int_{\mathbb{R}} \frac{1}{(2\pi)^{\frac{n}{2}} |\mathbf{D}|^{-\frac{1}{2}}} \exp\left(-\frac{1}{2} (\mathbf{R} - \mu \mathbf{1})^\top \mathbf{D}^{-1} (\mathbf{R} - \mu \mathbf{1})\right) \frac{1}{\sqrt{2\pi\sigma_\mu^2}} \exp\left(-\frac{1}{2\sigma_\mu^2} \mu^2\right) d\mu \\ &= \frac{1}{(2\pi)^{\frac{n}{2}+1} \sqrt{\sigma_\mu^2 \prod_{i=1}^n \sigma_i^2}} \exp\left(-\frac{1}{2} \sum_{i=1}^n \frac{R_i^2}{\sigma_i^2}\right) \end{aligned}$$

$$\begin{aligned}
& \times \int_{\mathbb{R}} \exp \left(-\frac{1}{2} \left(\frac{\mu^2}{\sigma_\mu^2} - 2\mu \sum_{i=1}^n \frac{R_i}{\sigma_i^2} + \mu^2 \sum_{i=1}^n \frac{1}{\sigma_i^2} \right) \right) d\mu \\
& = \left((2\pi)^n \left(1 + \sigma_\mu^2 \sum_{i=1}^n \frac{1}{\sigma_i^2} \right) \prod_{i=1}^n \sigma_i^2 \right)^{-\frac{1}{2}} \exp \left(\frac{1}{2} \left(\frac{\sigma_\mu^2 \left(\sum_{i=1}^n \frac{R_i}{\sigma_i^2} \right)^2}{1 + \sigma_\mu^2 \sum_{i=1}^n \frac{1}{\sigma_i^2}} - \sum_{i=1}^n \frac{R_i^2}{\sigma_i^2} \right) \right) \quad (13)
\end{aligned}$$

The likelihood ratio can now be computed by substituting Equation 13 into Equation 12 three times to arrive at:

$$\begin{aligned}
\frac{\mathbb{P}(\mathbf{R} \mid \mathfrak{F}_* \sigma^2, \gamma)}{\mathbb{P}(\mathbf{R} \mid \mathfrak{F}, \sigma^2, \gamma)} &= \sqrt{\frac{1 + \sigma_\mu^2 \sum_{i=1}^{n_\ell} \frac{1}{\sigma_i^2}}{\left(1 + \sigma_\mu^2 \sum_{i=1}^{n_{\ell_L}} \frac{1}{\sigma_i^2} \right) \left(1 + \sigma_\mu^2 \sum_{i=1}^{n_{\ell_R}} \frac{1}{\sigma_i^2} \right)}} \\
&\times \exp \left(\frac{\sigma_\mu^2}{2} \left(\frac{\left(\sum_{i=1}^{n_{\ell_L}} \frac{R_i}{\sigma_i^2} \right)^2}{1 + \sigma_\mu^2 \sum_{i=1}^{n_{\ell_L}} \frac{1}{\sigma_i^2}} + \frac{\left(\sum_{i=1}^{n_{\ell_R}} \frac{R_i}{\sigma_i^2} \right)^2}{1 + \sigma_\mu^2 \sum_{i=1}^{n_{\ell_R}} \frac{1}{\sigma_i^2}} - \frac{\left(\sum_{i=1}^{n_\ell} \frac{R_i}{\sigma_i^2} \right)^2}{1 + \sigma_\mu^2 \sum_{i=1}^{n_\ell} \frac{1}{\sigma_i^2}} \right) \right).
\end{aligned}$$

The Likelihood Ratio for the CHANGE proposal

The homoskedastic BART implementation of Kapelner and Bleich (2013) considered change proposals for singly internal nodes only (i.e., both daughter nodes must be terminal nodes). The likelihood ratio in this case simplifies to the likelihood of these two leaves before and after the change:

$$\frac{\mathbb{P}(R_{\ell^*,1}, \dots, R_{\ell^*,n_{\ell^*}} \mid \sigma^2, \gamma) \mathbb{P}(R_{r^*,1}, \dots, R_{r^*,n_{r^*}} \mid \sigma^2, \gamma)}{\mathbb{P}(R_{\ell,1}, \dots, R_{\ell,n_\ell} \mid \sigma^2, \gamma) \mathbb{P}(R_{r,1}, \dots, R_{r,n_r} \mid \sigma^2, \gamma)}. \quad (14)$$

The ℓ refers to the left terminal node and r refers to the terminal right node. The ℓ^* and the r^* denote these same two nodes in the proposal tree, i.e after the parent's split rule was changed.

The likelihood with μ margined out has been calculated in Equation 13 and we express it here with a convenient factorization:

$$\underbrace{\left((2\pi)^{n_\ell} \left(1 + \sigma_\mu^2 \sum_{i=1}^{n_\ell} \frac{1}{\sigma_i^2} \right) \prod_{i=1}^{n_\ell} \sigma_i^2 \right)^{-\frac{1}{2}}}_{A} \underbrace{\exp \left(-\frac{1}{2} \sum_{i=1}^{n_\ell} \frac{R_i^2}{\sigma_i^2} \right)}_{B} \underbrace{\exp \left(\frac{\sigma_\mu^2}{2} \left(\frac{\left(\sum_{i=1}^{n_\ell} \frac{R_i}{\sigma_i^2} \right)^2}{1 + \sigma_\mu^2 \sum_{i=1}^{n_\ell} \frac{1}{\sigma_i^2}} \right) \right)}_{C}.$$

To find the ratio, we must substitute this expression into equation 14 four times. We begin by substituting only the term marked *A* above to arrive at

$$\frac{(2\pi)^{\frac{n_\ell}{2}} \sqrt{\left(\prod_{i=1}^{n_\ell} \sigma_i^2\right) \left(1 + \sigma_\mu^2 \sum_{i=1}^{n_\ell} \frac{1}{\sigma_i^2}\right)} (2\pi)^{\frac{n_r}{2}} \sqrt{\left(\prod_{i=1}^{n_r} \sigma_i^2\right) \left(1 + \sigma_\mu^2 \sum_{i=1}^{n_r} \frac{1}{\sigma_i^2}\right)}}{(2\pi)^{\frac{n_{\ell^*}}{2}} \sqrt{\left(\prod_{i=1}^{n_{\ell^*}} \sigma_i^2\right) \left(1 + \sigma_\mu^2 \sum_{i=1}^{n_{\ell^*}} \frac{1}{\sigma_i^2}\right)} (2\pi)^{\frac{n_{r^*}}{2}} \sqrt{\left(\prod_{i=1}^{n_{r^*}} \sigma_i^2\right) \left(1 + \sigma_\mu^2 \sum_{i=1}^{n_{r^*}} \frac{1}{\sigma_i^2}\right)}}.$$

Of course $n_{\ell^*} + n_{r^*} = n_\ell + n_r = n$ and the σ^2 's are the same since they aren't drawn until later on in the Gibbs sampling scheme. Thus, the above reduces to

$$\sqrt{\frac{\left(1 + \sigma_\mu^2 \sum_{i=1}^{n_\ell} \frac{1}{\sigma_i^2}\right) \left(1 + \sigma_\mu^2 \sum_{i=1}^{n_r} \frac{1}{\sigma_i^2}\right)}{\left(1 + \sigma_\mu^2 \sum_{i=1}^{n_{\ell^*}} \frac{1}{\sigma_i^2}\right) \left(1 + \sigma_\mu^2 \sum_{i=1}^{n_{r^*}} \frac{1}{\sigma_i^2}\right)}}. \quad (15)$$

In term *B*, upon making all four substitutions, it is clear all the parents' observations must be summed in the numerator as well as the denominator. Thus, this term cancels.

Now we examine term *C*. Due to the exponentiation, multiplication becomes addition and division becomes subtraction and all four substitutions yield

$$\exp \left(\frac{\sigma_\mu^2}{2} \left(\frac{\left(\sum_{i=1}^{n_{\ell^*}} \frac{R_i}{\sigma_i^2}\right)^2}{1 + \sigma_\mu^2 \sum_{i=1}^{n_{\ell^*}} \frac{1}{\sigma_i^2}} + \frac{\left(\sum_{i=1}^{n_{r^*}} \frac{R_i}{\sigma_i^2}\right)^2}{1 + \sigma_\mu^2 \sum_{i=1}^{n_{r^*}} \frac{1}{\sigma_i^2}} - \frac{\left(\sum_{i=1}^{n_\ell} \frac{R_i}{\sigma_i^2}\right)^2}{1 + \sigma_\mu^2 \sum_{i=1}^{n_\ell} \frac{1}{\sigma_i^2}} - \frac{\left(\sum_{i=1}^{n_r} \frac{R_i}{\sigma_i^2}\right)^2}{1 + \sigma_\mu^2 \sum_{i=1}^{n_r} \frac{1}{\sigma_i^2}} \right) \right). \quad (16)$$

Multiplying terms 15 and 16 yields the likelihood ratio for the **CHANGE** proposal.